

Parallel Unsteady Overset Mesh Methodology for Adaptive and Moving Grids with Multiple Solvers

Matthew W. Floros*
Army Research Laboratory
Hampton, Virginia

Jayanarayanan Sitaraman
National Institute of Aerospace
Hampton, Virginia

ABSTRACT

This paper describes a new domain connectivity module developed to support Chimera-based interfacing of different CFD solvers for performing time-dependent, adaptive, and moving body calculations of external aerodynamic flows. The capabilities of the domain connectivity module are provide a powerful tool for CFD analysis of morphing vehicles. The domain connectivity module coordinates the data transfer between different solvers applied in different parts of the computational domain — body fitted structured or unstructured to capture viscous near-wall effects, and Cartesian adaptive mesh refinement to capture effects away from the wall. The CFD solvers and the domain connectivity module are executed within a Python-based computational infrastructure. The domain connectivity module is fully parallel and performs all its operations (identification of holes and fringe points, donor cell searches and data interpolation) on the partitioned grid data. In addition, the connectivity procedures are completely automated using the implicit hole-cutting methodology such that no user intervention or explicit hole-map specification is necessary. The capabilities and performance of the package are presented for several test problems, including flow over a NACA 0015 wing, AGARD A2 slotted airfoil, hover simulation of scaled V-22 rotor, and a dynamic simulation of UH-60A rotor in forward flight. A modification to the procedure for selecting the best data from multiple overlapping grids is also presented.

1.0 INTRODUCTION

Morphing vehicles present several unique challenges for computational fluid dynamics analysis. By definition, a morphing vehicle changes shape with time. The shape changing may be a gross deformation of the entire vehicle or the motion may be only small control surface or support structure. In either case, a single grid that captures both the important flow features around the entire vehicle and the details of its geometry while in motion can be challenging. Such problems lend themselves to multiple, overset grids. In the first case, where large parts of the vehicle are moving, grids wrapped around these large parts must undergo significant deformation to accommodate the shape change. In the latter case, only a small deformation may be required, but the purpose of the deformation may be to take advantage of unique flow physics that require localized refinement or even a specific solution methodology that is not necessary for the rest of the vehicle. The ability to combine heterogeneous grid types and solvers for an unsteady, moving body simulation is of great utility.

Traditional CFD codes are often written to support a single gridding and solution paradigm. Grids fall under three main classifications: Cartesian (structured or unstructured), structured-curvilinear (body-fitted) or unstructured (tetrahedral, hexahedral or prismatic). Each meshing paradigm has specific advantages and disadvantages. For example, Cartesian grids are easy to generate, to adapt, and to extend to higher-order spatial accuracy, but they are not well-suited for resolving boundary layers around complex geometries. Structured curvilinear grids work well for resolving boundary layers, but the grid generation process for complex geometries remains tedious and requires considerable user expertise. General unstructured grids are well-suited for complex geometries and are relatively easy to generate, but their spatial accuracy is often limited to second-order, and the associated data structures are less computationally efficient than their structured-grid counterparts.

Thus, while a single gridding paradigm brings certain advantages in some portions of the flow field, it also imposes an undue burden on others. A computational platform that supports multiple mesh paradigms provides the potential for optimizing the gridding strategy on a local basis. However, integrating different meshing paradigms into a single large monolithic code is

*Presenting author, matt.floros@us.army.mil, NASA Langley Research Center MS 266, 6C East Taylor Street, Hampton, VA 23681

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2010	2. REPORT TYPE		3. DATES COVERED 00-00-2010 to 00-00-2010		
4. TITLE AND SUBTITLE Parallel Unsteady Overset Mesh Methodology for Adaptive and Moving Grids with Multiple Solvers			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Army Research Laboratory, Hampton, VA, , ,			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) NATO/RTO			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Public Release	18. NUMBER OF PAGES 23	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

complex and usually relegates at least one of the models to be less accurate, less optimized, or less flexible than the original standalone solver.

This paper describes an approach to mitigate these issues by using multiple-mesh strategy that is implemented through the use of *multiple* CFD codes, each optimized for a particular mesh type. It was developed for analysis of rotorcraft, but features several enabling technologies for morphing vehicles and other types of unsteady, moving-body problems. In addition to supporting heterogeneous grids and solvers, the software was developed to minimize the analysis burden on the user. While an interface procedure is required for each specific solver, this interface must only be written once. Once the interface exists, the domain connectivity software retrieves the information it needs directly from the flow solvers and the grids themselves, so the human analyst does not need to provide input specific to each problem.

For rotorcraft analysis, the approach is to apply unstructured or body-fitted curvilinear grids near the body surface to capture complex geometry and viscous boundary layers. A short distance from the body surface, the flow is determined by a high-order block-structured adaptive Cartesian solver that adapts time-dependently to capture wake effects. Previous work [1, 2] has addressed the development of a Python-based multi-solver infrastructure using both unstructured (NSU3D [3]) and structured (UMTURNS [4]) near-body solvers and an adaptive Cartesian off-body solver (SAMARC). Using this high-level framework, discrete codes can be coupled together to obtain a contiguous solution over the entire computational domain. The Python code merely manages data pointers to pass information between the solvers and the domain connectivity package and adds negligible overhead to the analysis. Figure 1 shows an example calculation of flow over a sphere using this approach for which both the viscous boundary layer and the shed vorticity are accurately captured.




Figure 1: Unsteady flow over a sphere at $Re=1000$ using the multiple-solver approach. The NSU3D unstructured solver is used near the body surface, the high-order Cartesian AMR solver SAMARC is used in the field, and data is interpolated between the solvers using Chimera-based interpolation [1].

A critical aspect of the multi-mesh/multi-solver approach is the need for data exchange between the different meshes, which is facilitated in this work using the well-established Chimera-based overset procedure. As shown in Fig 2, the near-body and off-body meshes are constructed to be overlapping and the fringe data are interpolated using a domain connectivity module. Specifically, the domain connectivity procedures involve the evaluation of inter-grid boundary data points (points that receive data), donor cells (points that provide data), hole points (points that do not need to be solved) and interpolation weights. The focus of the present paper is on the development of a new domain connectivity module to support the overset multiple-mesh paradigm efficiently for large-scale unsteady computations and its applicability to unsteady moving body problems such as

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

morphing vehicles.

Several domain connectivity approaches have been investigated in the past by various research groups. The prominent among them are PEGASUS5 [5], OVERFLOW-DCF [6, 7], SUGGAR/DiRTlib [8], CHIMPS [9], BEGGAR [10], FAS-TRAN [11], and Overture [12]. Prior to the inception of the current effort, an evaluation of these packages showed that all of them had certain deficiencies that made them non-ideal for the multi-solver paradigm. For example, PEGASUS5 and SUGGAR, while being robust and validated, are intended to be integrated with a single solver as opposed to interfacing multiple solvers, and the domain connectivity operation is not parallel, which limits their applicability to large-scale moving body problems. OVERFLOW-DCF proved to be an efficient parallel solver for moving body problems, but it is tightly integrated within the OVERFLOW [7] code, making its applicability to a modular multiple solver paradigm difficult. Also, there is no support currently available in the module for unstructured meshes. The CHIMPS package was quite modular with an excellent API and had support for parallel execution in moving body environments, but it suffered from lack of support for automated hole-cutting and inefficiency in search procedures. All of the above packages required user input in some form or other for specifying hole-regions (i.e. regions of grid where flow solutions are not performed) and typically performed explicit hole cutting.




Figure 2: Overset near/off-body gridding. Unstructured or curvilinear grids to capture geometric features and boundary layer near body surface, adaptive block-structured Cartesian grids to capture far-field flow features.

The explicit hole-cutting methodology is known to be prone to several failure modes for moving body problems. An alternate approach is known as “implicit hole-cutting” in which holes and fringe points are determined as part of a donor search process rather than *a priori* by the analyst. This concept was initially investigated by PEGASUS5 [5] researchers and eventually deemed too inefficient for dynamic problems. However, more recently a research code, NAVAIR-IHC [13], was able to show efficient usage of this methodology for application to moving body problems. The NAVAIR-IHC code also was evaluated for potential application within the multi-solver paradigm. However, lack of parallelism, implementations that are specific to a structured grid topology, and failure modes for concave grid topologies made it unsuitable for the general application in the multi-solver paradigm with partitioned grid data.

For seamless usage in the computational infrastructure, the Domain Connectivity software should be modular, parallel, fully-automated, and efficient for unsteady moving body problems[14]. The primary subject of this paper concerns the development and application of such a methodology in the multi-solver context. In addition, this paper demonstrates the capabilities of the multi-solver paradigm in modeling complex problems with optimal grid distribution.

The new package developed as part of this work is termed PUNDIT, which is an acronym for Parallel Unsteady Domain Information Transfer. The objective of the PUNDIT developers is to create a package which is modular, parallel, and supports efficient automated domain connectivity operations for all participant grid types (unstructured, structured curvilinear and adaptive Cartesian). To minimize user inputs, the implicit hole-cutting methodology for identification of hole and fringe regions is used. In addition, PUNDIT operates in the same computational infrastructure as the participant codes and uses data pointers for grids and solution variables directly, thereby promoting ease of interfacing as well as reducing the memory footprint.

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

2.0 Methodology

The coupling of the near-body solver, off-body grid manager, off-body solver, and domain connectivity module are accomplished through a Python-based infrastructure with emphasis on preserving the modularity of the participating solvers. In addition to the advantages in efficiency and ease in code development, coupling existing mature simulation codes through a common high-level infrastructure provides a natural way to reduce the complexity of the coupling task and to leverage the large amount of verification, validation, and user experience that typically go into the development of each separate model. The infrastructure discussed here currently couples the following codes through a Python infrastructure: (a) the parallel NSU3D code [3] for the unstructured near-body solver, (b) the parallel UMTURNS code as the structured near-body solver, (c) the SAMRAI framework [15] for the off-body Cartesian grid generation, adaptation, and parallel communication, (d) the serial high-order ARC3DC code for solution on the Cartesian blocks, and (e) the domain connectivity module (PUNDIT).

Brief descriptions of the implementations of the first four components are outlined below. More detailed descriptions of the implementations as well as performance statistics and validation are described in more detail in Refs [1, 2]. This paper is devoted to detailed descriptions of algorithms and implementation of the domain connectivity module (PUNDIT). The infrastructure which integrates the simulation capabilities of all the modules mentioned above is called HELIOS (Helicopter Overset Simulations) and is developed under the DoD HPC HI-ARMS program. Because of the modular development approach being taken, the individual components of HELIOS, such as the solvers and the domain connectivity package, can be used independently for moving body problems other than rotorcraft.

2.1 Python Infrastructure

Python-based computational frameworks have been developed previously by several researchers [16] as a means of coupling together existing legacy codes or modules. Such a framework has a number of advantages over a traditional monolithic code structure: (1) it is easier to incorporate well-tested and validated legacy codes rather than to build the capabilities into an entirely new code, (2) there is less code complexity in the infrastructure itself, so maintenance and modification costs are less, and (3) it is easier to test and optimize the performance of each module separately, often yielding better performance for the code as a whole. Essentially, Python enables the legacy solvers to execute independently of one another and reference each other's data without memory copies or file I/O. Further, the Python-wrapped code may be run in parallel using pyMPI or myMPI, with each of the solvers following its native parallel implementations. Further details of the Python infrastructure used here can be found in Wissink et al. [1] and Sitaraman et al. [2].

2.2 Flow Solver Modules

Well-established legacy codes are employed for all of the independent modules in this study. For the unstructured solver in the near-body region, we employ the NSU3D [3] code, which is an implicit node-centered Reynolds-Averaged Navier-Stokes (RANS) code capable of handling arbitrary unstructured mesh elements. In addition, we also use a curvilinear structured mesh solver for the near-body region, namely the UMTURNS code, which is also an implicit RANS solver. For the Cartesian grids in the off-body region, we utilize a Cartesian grid derivative of the well-known ARC3D [17] code, referred to here as ARC3DC. The ARC3DC code employs third-order temporal discretizations using a multi-stage Runge-Kutta time-stepping framework and is capable of up to fifth-order accurate spatial discretizations. Further, the Cartesian grids in the off-body are automatically generated and managed for parallel execution by the SAMRAI infrastructure [18, 15]. As mentioned earlier, these codes or modules are combined together using a Python-based framework that orchestrates the execution of these modules and the associated data transfers.

2.3 Meshing Paradigm

The meshing paradigm consists of separate near-body and off-body grid systems. The near-body grid typically extends a short distance from the body, sufficient to contain the boundary layer. This grid can be a structured curvilinear grid or an unstructured tetrahedral or prismatic grid that has been extracted from a standard unstructured volume grid or generated directly from a surface triangulation using hyperbolic marching. The reason for using curvilinear or unstructured grids in the near-body region is to properly capture the geometry and viscous boundary layer effects, which are difficult or impossible to capture with Cartesian grids alone. We further note that either structured or unstructured grids will work equally well in the near-body region from the point of view of our infrastructure. Away from the body the near-body grid solution is interpolated onto a Cartesian

background mesh with the aid of the domain connectivity algorithm. This transition normally occurs at a distance wherein the sizing of the near-body grid cells is approximately commensurate with the sizing of the Cartesian mesh in the off-body region.




Figure 3: Block structured AMR grid composed of a hierarchy of nested levels of refinement. Each level contains uniformly-spaced logically-rectangular regions, defined over a global index space.

The off-body grid system consists of a hierarchy of nested refinement levels, generated from coarsest to finest. In the Structured Adaptive Mesh Refinement (SAMR) paradigm [19], the coarsest level defines the physical extents of the computational domain. Each finer level is formed by selecting cells on the coarser level and then clustering the marked cells together to form the regions that will constitute the new finer level. Solution-based refinement progresses as follows: physical quantities and gradients are computed at each Cartesian cell using the latest available solution and those cells that hold values deemed to require refinement are marked. The marked cells are then clustered to form the new set of patches that constitute the next finer level. The process is repeated at each new grid level until the geometry and solution features are adequately resolved. We note that this entire procedure is automated within the software and no user intervention is required. The procedure of adaptive mesh refinement is graphically illustrated in Fig. 3.

An example of the overset near-body/off-body meshing strategy is given in Fig. 2, which shows the meshes for flow computations over a helicopter fuselage. Here, the mixed-element unstructured near-body grid envelops the fuselage, while a multi-level Cartesian off-body grid extends from the outer boundary of the near-body grid to the far-field boundary. The two sets of meshes overlap in the so-called fringe region, where data are exchanged between the grids.

2.4 Overset Methodology

The overall solution procedure for the overset meshes is as follows. At each iteration step, the solution of the fluid equations in each mesh is obtained independent of each other with the solution in the fringe region being specified by interpolation from the overlapping “donor” mesh as Dirichlet boundary conditions. At the end of the iteration, the fringe data is exchanged between the solvers so that the evolution of the global solution is faithfully represented in the overset methodology. Further, if one or more of the meshes is moving or changing, the fringe regions and the interpolation weights are recalculated at the beginning of the time step. This procedure is repeated for each iteration step until solution convergence is attained in both the near- and off-body meshes.

2.5 Domain Connectivity Module (PUNDIT)

For each solver, the solution at the fringe cells is obtained by interpolation from the overlapping “donor” mesh. The domain connectivity module (PUNDIT) is responsible for determining the appropriate interpolation weights for each fringe point. Further, in the case of multiple overlapping meshes, PUNDIT must also identify one mesh as the donor for each fringe point in each mesh. For static meshes, these operations are done once, at the beginning of the computation, while, for the more general case of moving or adapting meshes, the determinations of donors and weights has to be repeated within the time-stepping or iteration loop.

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

PUNDIT adopts the implicit hole cutting procedure followed by NAVAIR-IHC [13]. The core idea of this approach is to retain the grids with the finest resolution at any location in space as part of the computational domain and interpolate data at all coarser grids in this region from the solution on the fine grid. This results in the automatic generation of optimal holes without any user specification as in the case of explicit hole-cutting. Moreover, the implicit hole-cutting procedure produces an arbitrary number of fringe points based on mesh density compared to traditional methods which use a fixed fringe width (usually single or double). The critical parameter that quantifies the quality of a grid cell or node is termed as “resolution capacity”. PUNDIT nominally uses the cell volume as the resolution capacity for a grid cell and the average of cell volumes of all associated grid cells as resolution capacity for a grid node. A modification to this procedure to account for wall proximity is described later. In addition, PUNDIT separates the near-body to near-body and near-body to off-body domain connectivity procedures to facilitate automatic off-body grid generation and to improve efficiency (Meakin et al. [14]).

Following are the steps followed by PUNDIT to determine the domain connectivity information in a parallel environment. Partitioning of grid and solution data is assumed to be performed using an appropriate solver-based load balancing scheme before these steps are executed.

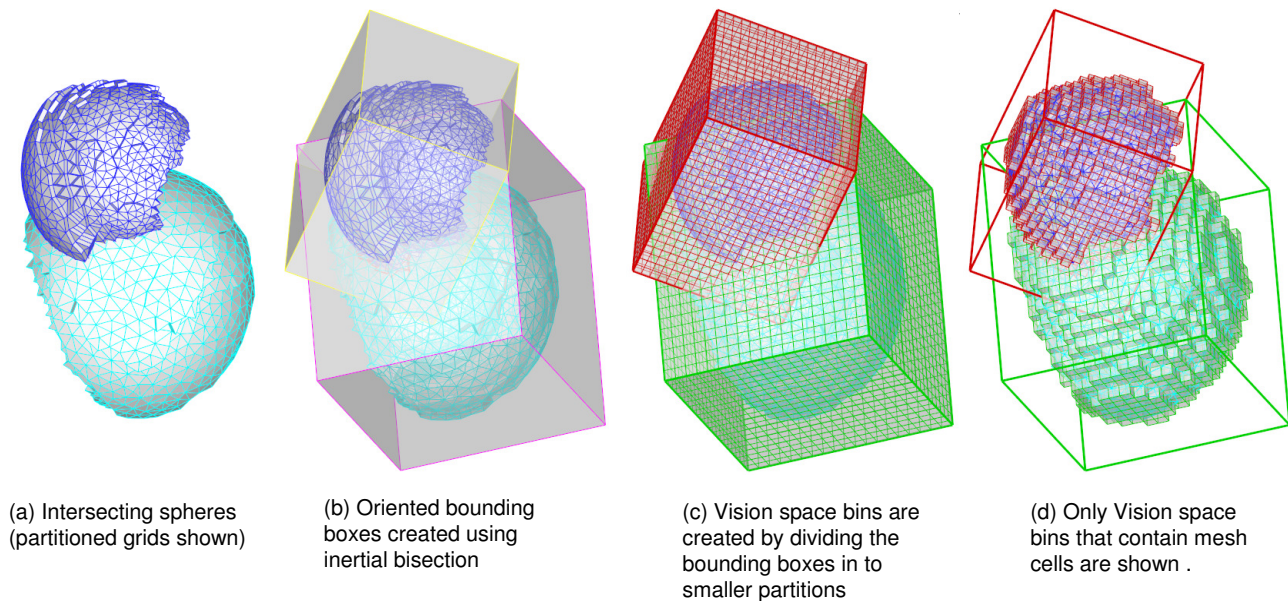


Figure 4: Oriented bounding boxes and vision space bins that are created during preprocessing to accelerate donor search processes.

1. *Registration*: On each processor, the flow solvers register grid and solution data pointers for every grid block (both near body and off-body) with PUNDIT.
2. *Profiling*: PUNDIT profiles each of the grid blocks and forms meta-data representations to facilitate faster donor search operations. The main procedures that are executed in this step are:
 - *Minimal bounding box* computation: Oriented bounding boxes are constructed instead of axis-aligned bounding boxes to minimize the search space (Figure 4(b)).
 - *Division of minimal bounding box* in to *vision space bins*: The vision space bins are smaller Cartesian boxes within the bounding box. The size of the vision space bins is determined by dividing the volume of the bounding box by the total number of cells contained in it (Fig 4(c) and Fig 4(d)).
 - *Generation of a cumulative fill-table*: The cumulative fill table is a bin-wise reordering of the cell indices that are contained in each vision space bin. They facilitate fast identification of all cells contained within each vision space bin.

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

- Estimation of element *resolution capacity*: Element resolution capacity is the grid quality measure for each cell and each grid node. Nominally, cell volumes are used as resolution capacity for each grid cell and average of all the associated cell volumes is used as the resolution capacity for each grid node. A method to modify resolution capacity to retain cells near a solid wall in the grid containing the wall is described in section 4.0.
3. *Near-body to near-body domain connectivity*: In this step the domain connectivity operations are performed between all near-body blocks on all processors. The sequence of the operations is as follows: The bounding boxes constructed in each processor in the previous step are gathered in every processor. Each processor checks for intersection of bounding boxes of its grid blocks with the global set of bounding boxes, which aids in identifying candidate donor grid blocks. The bounding box intersection checks are optimized such that detection of intersection as well as identification of vision space bins containing potential receiver grid points are simultaneously processed. The identification of vision space bins accelerates the identification of potential receiver grid points as only those nodes which belong in these vision space bins are checked for containment in the donor bounding box.

Following the intersection checks a communication packet is set up and exchanged between all the processors. The communication package consists of a list of coordinates of all potential receiver grid points and their resolution capacity organized such that the donor search can be directed to the appropriate grid block in the candidate processor. Upon completion of this communication, each grid block in each processor will form a list of points for which it needs to locate donor cells. The donor search is then conducted following the algorithm outlined below. Once a list of donors is obtained, it is communicated back to the processor that requested the donor search, where an evaluation of potential donor cells which were found is performed, such that the donor of the best resolution capacity is chosen for every receiver grid point. Indices of all the potential donors that are not acceptable are communicated back to the appropriate donor processor such that it can update its donor list. Additionally a quality check is performed to make sure that no donor cells have a receiver point as its vertex. If such a cell is located it is deleted from the donor list and this information is communicated to the receiver processor which adjusts its receiver list accordingly as well. This process rigorously establishes donor quality since all donors will be composed of only nodes which are being solved and are not themselves receiver points. The final product of this step is a communication table consisting of a list of donors and receivers in each processor which is synchronized for data interpolation.

Figure 5 shows a graphical illustration and flow chart of the donor search algorithm outlined below. The donor search algorithm proceeds as follows: For any potential receiver point, a localization is performed by locating the vision space bin that contains it. This is a trivial operation since the vision space bins follow a Cartesian structure within the oriented bounding box. Subsequently a spiral search is performed beginning at the vision space bin until a bin with at least one grid cell is located. This grid cell is chosen as the seed for initiating the so called “stencil walk process”. A line is created by connecting a point inside a potential donor cell (centroid for the initial cell and face intersection point for subsequent cells) with the receiver point. A check is made to determine if this edge intersects any of the faces of the cell. If an intersection is located, the cell which forms the neighbor at that face is chosen as the next potential donor cell. This procedure continues until a cell with no intersections is located which would be the donor cell for the given point. For a face which shows intersection but has no neighbor (i.e. a boundary face), a check is conducted using the spiral search technique to check if the edge intersects any other boundary faces in the neighborhood. If an intersection is found then stencil walking proceeds to the cell which owns that boundary face. This procedure addresses complex grid boundaries such as those found in partitioned unstructured grids or thin trailing edges. If a cell is not found, the receiver point is pronounced to have no donor cell in the current grid block. In addition, if the final boundary face that the line intersected is a wall boundary face the receiver point will be tagged as a hole point (i.e. it is inside the solid body). Once a point is tagged as a hole by a near body grid, it will no longer receive data even if a suitable donor is found by another grid. This ensures that the volume within the body is blanked for the nodes in all grids which overlap within the body.

4. *Off-body grid generation*: The adaptive Cartesian mesh generation is dictated by the resolution of the near body meshes at their boundaries. In order to achieve this, a list is generated which is composed of all outer boundary points of near body meshes across processors that did not find a near-body donor in the previous step. This list is communicated to the off-body grid manager (SAMRAI) which automatically generates/adjusts nested Cartesian meshes such that the resolution is commensurate at the near-body boundaries.
5. *Off-body to Near-body connectivity*: First valid donors (i.e those with better resolution capacity) for all points in near-body grid blocks are searched in the off-body Cartesian blocks. This search process is extremely efficient because of the isotropic nature of the Cartesian blocks and the global grid indexing followed by the off-body grid manager. A donor Cartesian cell can be located in just a single step. In addition, this process identifies all the Cartesian blocks which might

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

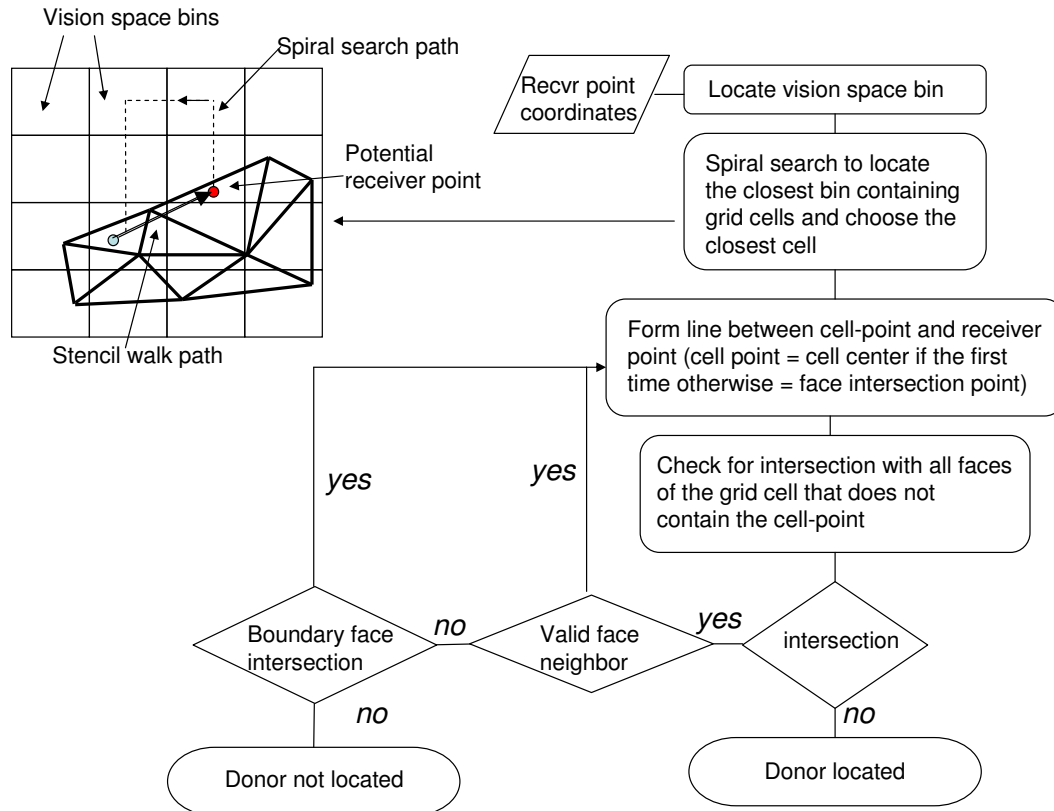


Figure 5: Flowchart of donor finding algorithm that uses vision space bin based localization and stencil-walk based search.

contain potential receiver points. Following this, valid donors in the near-body grids are searched for all potential off-body grid points using the same process as that was outlined for the near-body to near-body domain connectivity. Once the donors and receiver points are located for both the off-body and near-body grids, a communication pattern is followed which updates the communication tables that synchronize the donor and receiver lists in each processor. Donor quality is again guaranteed by removing low quality donors and their corresponding receivers from the communication tables. Once the communication tables are finalized, the interpolation weights for each grid node of each donor cell in the donor list are computed. Tri-linear basis functions are used for data interpolation and the interpolation weights are computed using a Newton-Raphson procedure.

6. *Data interpolation:* Data is interpolated in a three step process. The first step consists of populating a communication buffer using the solution data and the donor list from the communication tables. These data buffers are exchanged between processors through interprocessor communication in the second step. In the third and final step the solution data is updated in each processor based on the receiver list from the communication tables. In contrast to many of the existing domain connectivity solvers, PUNDIT minimizes the communication overhead by maintaining the interpolation weights in the host processor. The client processor only receives solution data and is capable of assigning them appropriately because of the synchronization in communication tables. Moreover, interpolation weights are evaluated only for the final list of donors, minimizing the amount of floating point operations. In order to support interfacing of solvers which follow different non-dimensionalizations for flow variables, PUNDIT maintains non-dimensionalizing factors for each flow variable from each participant solver. The data buffers that are exchanged are in the dimensional form and appropriate non-dimensionalization is applied at the update step based on the factors provided by each participant code. It is worth noting that the implementation of data interpolation is quite general and does not impose any restriction on the type or number of flow variables.

3.0 Software Verification and Application Results

We present results using PUNDIT within the flow solver package HELIOS for several examples. The results section is organized as follows: first, test cases for interpolation accuracy and scalability are presented. Then, application results are presented which demonstrate the use of the software for adapting and moving body meshes. The first result is a NACA0015 fixed wing test case which demonstrates the advantages of adaptive grids and higher-order methods for wake capturing. The next set of results illustrates the capability of PUNDIT to perform domain connectivity operations which automate hole cutting and generate optimal overlaps. The domain connectivity solutions are shown for the AGARD A2 test case (NHLP-2D slotted airfoil) in steady and unsteady motion. Following that we present results for rotorcraft simulations for the 1/4th scaled V-22 and full scale UH-60A rotors. Additional application results are presented in Ref. [20].

3.1 Interpolation Accuracy

All the polyhedra that are common in unstructured mixed-element meshes (tetrahedra, pyramids, prisms and hexahedra) are supported in PUNDIT. Tri-linear basis functions are used to perform data interpolation inside each of these polyhedra. The accuracy of interpolation was verified by constructing unstructured meshes in a unit cube of incremental resolutions. PUNDIT is used then to perform search and interpolation on a random set of points distributed in the unit cube using the data at the grid nodes (which are prescribed using chosen test functions) of the unstructured meshes. The interpolated values are compared with exact test function values to estimate the interpolation error. Figure 6 shows the variation of interpolation error with improving resolution. For linear test functions, the tri-linear interpolation gives exact solutions within machine precision. For a non-linear test functions interpolations in all types of polyhedra show 2nd order accuracy. Interpolation accuracy is also verified to be 2nd order for a test function that uses the Lamb-vortex solution.




Figure 6: Variation of interpolation errors for polyhedra shown in subplot (d): subplot(a) shows error variation for a linear test function, subplot(b) shows error variation for a non-linear test function and subplot(c) shows error variation for the Lamb-vortex solution.

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

3.2 Scalability Studies for UH-60A and TRAM Rotors

The performance of the domain connectivity module is studied by varying the number of processors used for program execution. Note that partitioning of the grid data is handled by flow solvers and PUNDIT operates on this partitioned grid data. Figure 7 shows the speedup in wall-clock time with increasing number of processors. Three sets of grid systems are studied to calibrate scalability. The TRAM grid system consists of 0.8 million near body nodes and 5.9 million off body nodes. The coarse UH-60A grid system is composed of 1 million near-body nodes and 11 million off-body nodes. The fine UH-60A grid system contains 2 million near-body nodes and 30 million off-body nodes. More details on these test cases are described in the application results section (3.6 and 3.5). Good linear scalability was observed for all three cases up to 12 processors. Beyond that the scalability drops off depending on grid size. Larger mesh systems such as the fine grid for UH-60A begin to drop off at about 32 processors, while the TRAM rotor starts to drop much earlier at about 12 processors. The reason for scalability drop-off is the load-imbalance created by non-optimality of solver based partitioning for domain connectivity procedures. Redistribution of the computational work for searching donors evenly after an initial evaluation of donor processors could be an approach to mitigate this problem (For e.g. the “rendezvous” approach described in Plimpton et al [21]).




Figure 7: Scalability of PUNDIT with increasing number of processors for TRAM and UH-60A data sets.

3.3 Flow Over NACA0015 Wing

The multi-code approach is used to study the flow-field over a NACA 0015 rectangular wing. The primary objective of this computation is to validate the multi-code simulation capability for a high-Reynolds number turbulent external flow problem. The experimental data used for validation are those measured by McAlister et al. [22]. The wing geometry is a rectangular planform with a square-tip and has an aspect ratio of 3.3. The test point for which validation results are presented is at the following operating condition: $M=0.1235$, $Re=1.5e6$, $\alpha=12^\circ$. For this case, computations are performed in the near-body region using the UMTURNS flow solver using the structured curvilinear meshes. The wing is modeled in the full-span configuration. The Spalart-Allmaras turbulence model is used for RANS closure.

Figure 8 shows the domain connectivity solution obtained for the NACA0015 grids. Sections of the grid system in the X-Y plane and X-Z plane are shown. The implicit domain connectivity algorithm employed not only produces optimal overlap but also ensures commensurate grid spacing of off-body and near-body computational domains at regions of overlap. The pressure distributions obtained from the predictions are compared with measured experimental data in Figure 9. Fair agreement can be noted between the measurement and prediction at all the span locations, although the leading edge suction peak is consistently under-predicted, This may be attributed to wind-tunnel wall effects which are not modeled in the analysis presently.




Figure 8: Domain connectivity solutions with adaptive Cartesian grids for steady flow over a NACA0015 wing at $M=0.1235$, $\alpha=12$ deg and $Re=1.5e6$ using HELIOS.

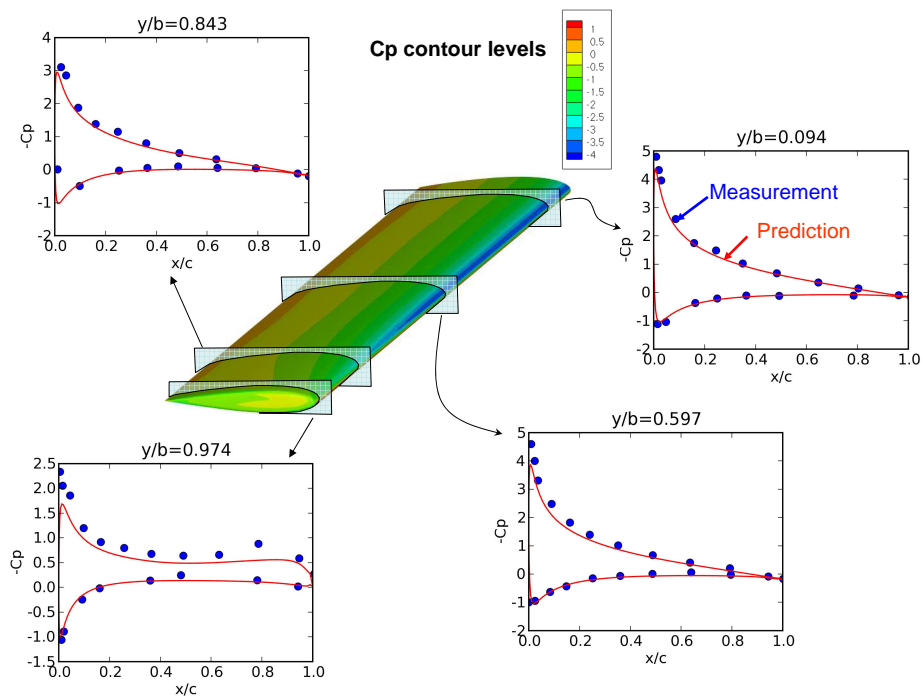


Figure 9: Prediction of pressure distributions for steady flow over NACA0015 wing at $M=0.1235$, $\alpha=12$ deg and $Re=1.5e6$ using HELIOS.

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

Figure 10 shows the evolution of the vortex structures from the tip of the wing. The flow field is dominated by the presence of the vortex sheet which rolls up into a strong tip vortex quite close to the wing. This tip vortex remains coherent up to about 15 chords behind the wing and then begins to break down in to smaller structures. Vorticity magnitude is used as the mesh adaptation criterion for the off-body solver. The meshes represented in the figure show four levels of adaption. It can be observed that the mesh system is able to track the vorticity and automatically adapt to regions of high vorticity (around the area of tip vortices). The figure also shows comparison of swirl velocities predicted at the tip-vortex locations with measured experimental data. The peak magnitudes are under predicted by about 20 percent which is consistent with the under-prediction of the leading edge suction peaks. However, the vortex shows minimal diffusion because of numerical dissipation. The use of a 5th-order accurate numerical scheme (6th-order central difference with 5th-order dissipation terms) in the off-body solver in conjunction with adaptive Cartesian grids facilitates the improved resolution of the tip vortex structure. This test case illustrates how the domain connectivity module can be used with an adaptive grid solver to capture localized flow features in the solution.




Figure 10: Prediction of steady flow over NACA0015 wing at $M=0.1235$, $\alpha=12$ deg and $Re=1.5e6$ using HELIOS; Predicted tip-vortex trajectory and comparison of predicted swirl velocities with experimental data.

3.4 AGARD A2 Slotted Airfoil Test Case

This section illustrates the capabilities of PUNDIT for the AGARD A2 (NHLP 2D slotted airfoil) test case in static and dynamic conditions. Although topologically simple, this problem is challenging for computation of domain connectivity because of multiple overlapping meshes in close proximity. In particular, when the flap and slat are fully retracted, there is a very narrow gap between the moving parts and the stationary airfoil body. For this case, there are independent grids for the flap, slat, and

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

main airfoil body, plus a background grid that extends to approximately 10 chords around the airfoil. For this test case, the UMTURNS code was used to calculate the flow solution in all four grids. This slotted airfoil is one of the well established cases for validation of CFD codes [23, 24].

The experimental data for this case are obtained from the electronic supplement of the AGARD 303 report [25, 26]. Only static data is available in the supplement. The surface pressure distributions obtained by the computations and their comparison with the experimental data for the 4 deg angle of attack case are shown in Figure 11. Excellent agreement can be noted in the suction peaks and chordwise variation of pressure coefficients for all three elements of the slotted airfoil. This validation study provides confidence in the accuracy of the domain connectivity module.




Figure 11: Prediction of steady 2-D flow over NHLP-2D slotted airfoil at $M=0.1937$, $\alpha=4$ deg and $Re=3.52e6$ using HELIOS; Comparison of surface pressure distributions to experimental data

Figure 12(a) shows the original overlapping meshes that describe the geometry of the slotted airfoil problem. Figures 12(b) and 12(c) show the detail of the leading and trailing edge section which contain regions which have three overlapping meshes. Finally, Figures 12(d) and 12(e) show the results after performing the domain connectivity operation. In these figures only the solver points are shown. It can be noticed that at the solver point boundaries the grid resolutions of all overlapping meshes are commensurate with each other which improves the quality of data interpolation. In addition the amount of overlap is optimized, which in turn minimizes the amount of double solving and redundancy in the flow solution.

Figure 13 shows the contours of stream-wise momentum near the leading edge slat and trailing edge flap. The continuity of the contours across the mesh boundaries illustrates the quality of data interpolation achieved by maintaining optimal connectivity. Data interpolation between mesh cells and nodes of differing resolution capacity often causes large, non-physical oscillations in contours of flow variables at grid interfaces.

Following the static computations, calculations were performed under dynamic conditions to illustrate the utility of the domain connectivity module for moving body and morphing vehicle problems. The flap and slat were cyclically oscillated at a 1 Hz frequency for two cycles with 8000 time steps per cycle. Although validation data for the flow solution is not available for the dynamic case, the performance of the domain connectivity software can be qualitatively shown. This type of problem demonstrates the power of the implicit hole cutting procedure. For the computation, the flap and slat grids were translated and rotated incrementally at each time step and the domain connectivity re-calculated with the new grid coordinates. The airfoil body and background meshes were stationary. As the flap and slat approach the airfoil body, an increasing number of grid points become hole points and background points become solver points. When fully retracted, the gaps between the airfoil body and the control elements are very thin, but PUNDIT is able to perform the hole cutting reliably and automatically without

**PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE
AND MOVING GRIDS WITH MULTIPLE SOLVERS**




Figure 12: Domain connectivity operations for NHLP airfoil-flap-slat test case using PUNDIT.

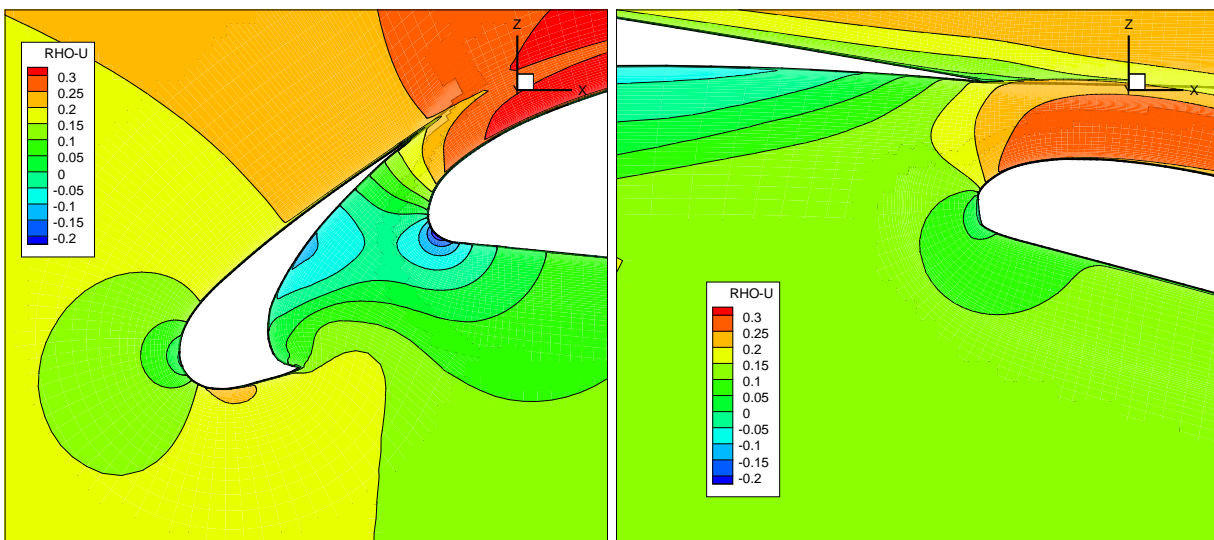


Figure 13: x -momentum contours of flow solutions for NHLP airfoil-flap-slat test case using PUNDIT and UMTURNS; Conditions are $M=0.197$, $\alpha=4$ deg, $Re=3.52e6$

any input from the user, see Figure 14. Explicit hole cutting for a complex shape with such a tight gap would be challenging and would require significant user expertise and effort.




Figure 14: AGARD A2 airfoil grids with flap and slat fully retracted after domain connectivity operation. The thick lines represent the boundaries of the grids before blanking.

For a sinusoidal oscillation, the velocity of the flap and slat are highest when they are 50% deployed. Figure 15 shows the streamwise momentum solution for the airfoil with the flap and slat in the same position while retracting and deploying. The large orange and red areas on the upper surface represent the flow accelerating over the airfoil. Because of the lag in the unsteady lift, the air is moving faster over the airfoil as the elements retract than as they deploy.

3.5 Hover Simulation of TRAM Rotor

Hover simulation for the TRAM (Tilt Rotor Aeroacoustics Model, a 1/4 scaled V-22) rotor was performed using the HELIOS infrastructure with PUNDIT for domain connectivity. Figure 16 illustrates the mesh system, connectivity solution, and flow solution obtained for this test case. This test case also shows the capability of PUNDIT to handle multiple unstructured meshes. The near-body solver used for this simulation is the unstructured NSU3D code.

The TRAM mesh system consists of three unstructured grid blocks generated by rotating and combining a single blade unstructured mesh. PUNDIT performs domain connectivity between these grid blocks first and determines the resolution and geometry requirements for the off-body meshes. The off-body mesh generation is performed subsequently such that the resolution of the grid cells near the outer boundaries of near-body meshes is commensurate with the off-body grid cells in that region. Consequently PUNDIT performs the domain connectivity solution between the off-body and near-body systems and determines the off-body and near-body computational domains. The off-body computational domains are shown overlapped with near-body meshes to illustrate the extent of overlap created. In addition the off-body domains are shown by themselves to illustrate the hole geometry created by the implicit hole-cutting procedure. Overall, it can be observed that PUNDIT is capable of producing optimal domain connectivity solutions for the mesh system used in this case, even at the center of the rotor where there are multiple overlap regions. It is worth noting that an explicit hole cutting methodology needs sufficient user intervention to make the appropriate hole cuttings for a mesh system such as this, especially at the overlap regions.

The flow solution obtained from the hover simulation, shown in Figure 17, illustrates the dynamic adaptation of the off-body grid system to the tip vortex structure. The solution slice shows contours of downwash momentum superimposed on the grid system. Large gradients of downwash indicate a larger presence of vorticity as observed in the tip regions where the mesh density is automatically increased. Note that domain connectivity is performed on the fly in this simulation after each off-body grid adaptation step.

3.6 Moving Mesh Computations for Rotor Aeromechanics

The last case shown is the moving mesh problem for the UH-60A helicopter in forward flight. The meshes used for this problem are shown Figure 18. Figure 18(a) shows the grid system with a curvilinear C-O type topology (near-body) overlapped by a


**PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE
AND MOVING GRIDS WITH MULTIPLE SOLVERS**

Figure 15: AGARD A2 airfoil solutions with the flap and slat in mid-motion. In the upper figure, the elements are retracting; in the lower figure, they are deploying.

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE
AND MOVING GRIDS WITH MULTIPLE SOLVERS




Figure 16: Mesh system and domain connectivity for 1/4 scaled V-22 (TRAM) rotor using HELIOS; Plots illustrate grid overlap and hole geometry.




Figure 17: Flow solutions for 1/4 scaled V-22 (TRAM) rotor using HELIOS; Plot illustrates dynamic mesh adaption to flow features.

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

Cartesian mesh (off-body). The flow solutions for the near body grids are performed by the UMTURNS code and those for the off-body grids are performed by the SAMARC code. Figures 18(b) and 18(c) show sections along the $z=0$ plane and $x = 0$ plane for the entire grid system illustrating the computational domains (solver points) in the near- and off-body regions. The problem includes grid deformation to account for blade elasticity as well as rigid body rotation of the near-body meshes around the rotor shaft axis. Prescribed elastic deformations which are obtained from previous CFD/CSD computations [27] are used for verification of the methodology. Since the near-body meshes are moving, the domain connectivity has to be performed at every time step. The domain connectivity procedures required about 10 percent of the total time taken for a time step. Computations were conducted on 32 processors, which is the scalability limit of PUNDIT for this problem.

Figure 18(d) shows the contours of the z -momentum (downwash) for the flow solution obtained. The vorticity shed from the rotor blades induces the downwash observed in the figure. During the advancing azimuth the aerodynamic loading towards the tip of the blades becomes negative causing them to produce much less downwash compared to the retreating side where the aerodynamic loading is positive. The downwash distribution is also an indicator of the vortex wake geometry shed from the rotor blade; large gradients indicate the presence of strong vortex structures. A comparison of aerodynamic loading (sectional lift and pitching moment at 86 percent radial station) is shown in Figure 18(e) which compare favorably with the measured flight test data [28].

4.0 Resolution Capacity Modification

As stated in section 2.5, the nominal resolution capacity that determines the solution quality at a given point in a given grid is based on the cell volume. For a cell, it is the volume of that cell, and for a point, it is the average of the volumes of the cells which share that point. In a purely mathematical sense, this provides the best solution by selecting closely spaced points over more sparsely spaced points. However, physically, the points with the closest spacing may not be the points which should be solved. Two particular instances were observed where the default selection process was not selecting the correct points. Both can be illustrated by the slotted airfoil case presented in section 3.4.

The first case is near a solid wall. For overset grids, only the solver assigned to the grid block which contains the wall has information that the wall is present. Processors solving background grid blocks do not have the wall information from the near body grids, and each near-body grid does not have information about walls in other near-body grids. Wall distance is important for turbulence modeling in particular, so it is important points near a solid wall are solved in the grid which contains the wall.

In the second case, a wall may not be involved but an “island” can form where one grid is locally more refined than another. Computationally, it may be important for solution points to be contiguous rather than being broken up by receiver points which get data from other processors. Localized areas of dense points are normally because of the proximity of a wall, so often the first issue is the cause of the second.

In order to address these issues, a wall correction scheme was investigated. A scaling factor ξ was applied to artificially promote the resolution capacity of cells near a solid wall. The scaling factor is identically 1.0 for grid blocks which do not contain any wall nodes and varies from 0 to 1 in grid blocks which do have wall nodes. This is performed by calculating ξ for all nodes, and then non-dimensionalizing by ξ_{\max} for the grid block so all values of ξ fall between 0 and 1.

The first method was a linear scale factor. As the resolution capacity is being calculated, the distance to the nearest wall node, x_w is calculated, and the scale factor was determined simply as

$$\xi = x_w \quad (1)$$

For this case, the ξ_{\max} is the distance of the boundary node farthest from a wall in the grid block. A failure mode of this approach is seen with the AGARD A2 airfoil grids. The grid for the main airfoil body extends far beyond the trailing edge compared to the grids for the flap and slat. Since the scale factor is relative the farthest point from the wall, the non-dimensionalized ξ is smaller for the main body than for the flap and slat for the same distance from the wall.

For the general case where grids can be heterogeneous and can contain complex shapes or protuberances, ξ should reflect the proximity to the wall relative to the local vicinity of the grid. To accomplish that, for each point, the shortest distance to an outer boundary x_o is also determined. With both the minimum distances to a wall and to an outer boundary, ξ was determined as

$$\xi = \left(\frac{x_w}{x_w + x_o} \right)^2 \quad (2)$$

Although the nearest wall point and the nearest outer point are not generally along the same line, this ensures that ξ approaches 0 near the wall and approaches 1 near the outer boundary. And for the case of the airfoil body mentioned above, the

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

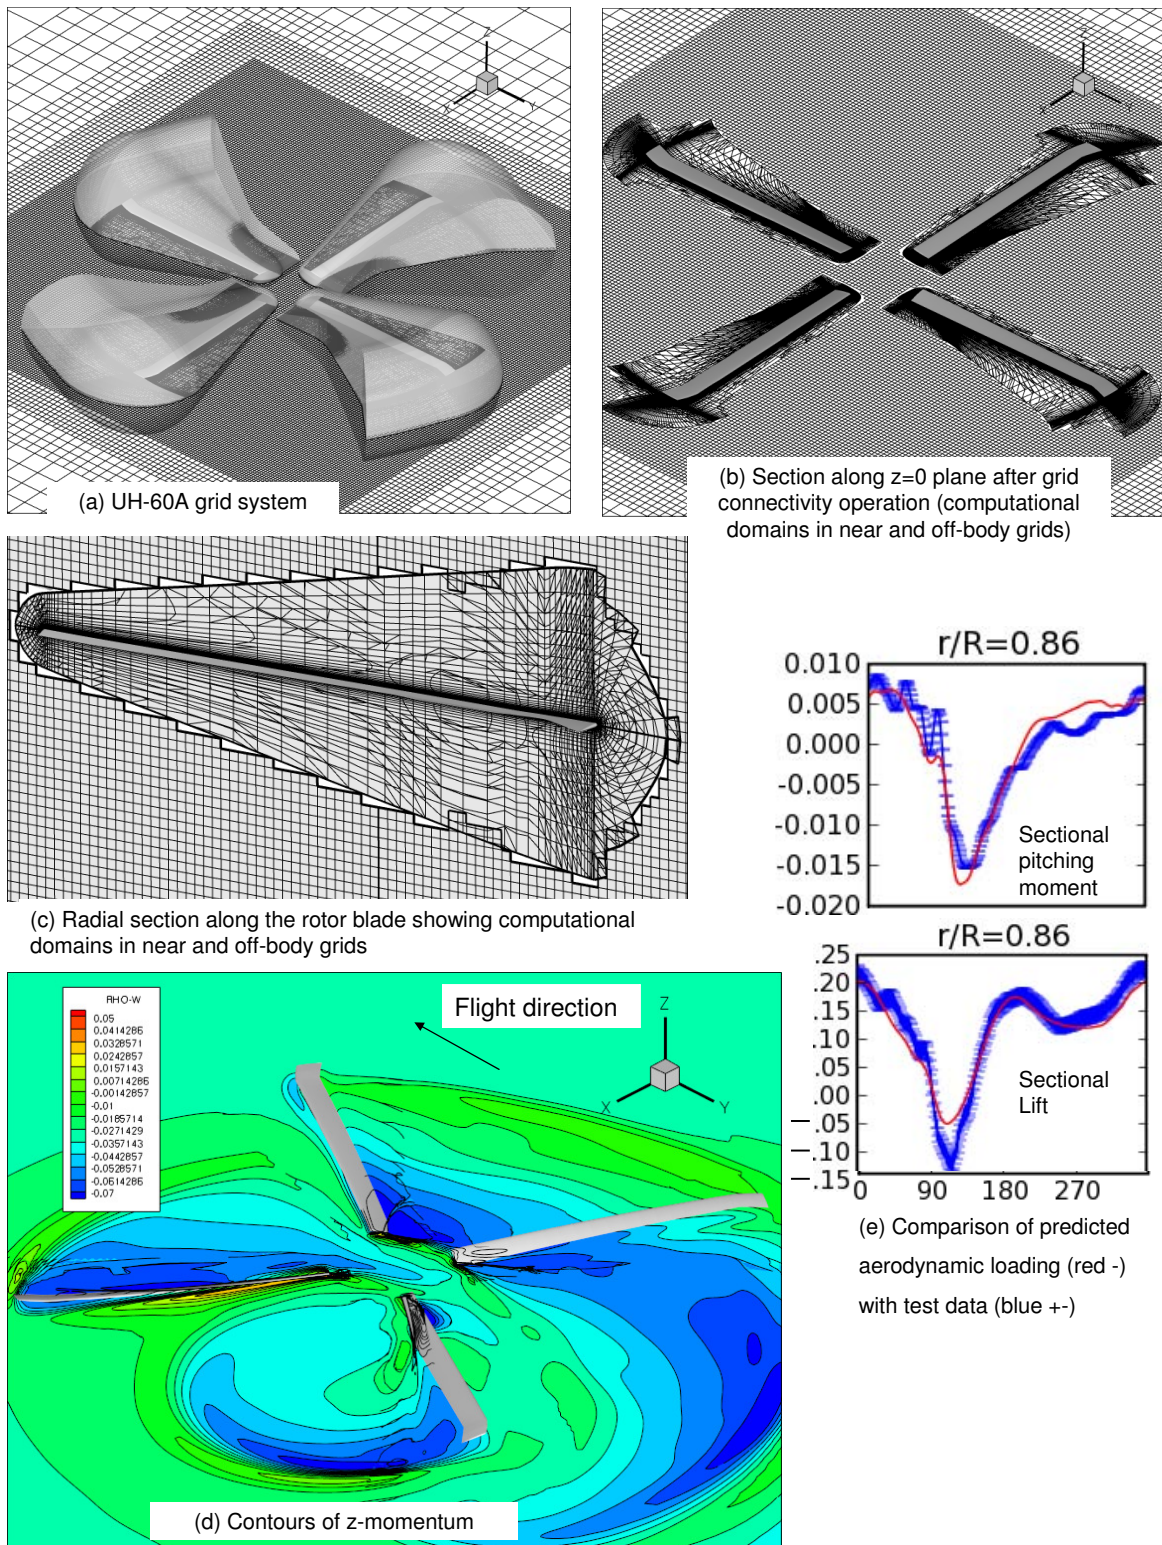


Figure 18: Mesh system, domain connectivity and flow solution for aeroelastic simulation of UH-60A rotor using HELIOS; Plots illustrate grid overlaps, hole geometry and comparison with measured flight test data.

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

resolution capacity of points near the body are not excessively modified because the grid extends so far beyond the trailing edge. For this investigation, the ξ values are still divided by ξ_{\max} even though it is close to unity. This ensures that the maximum scale factor for points at the outer boundary of grid blocks is exactly 1 to match that of background grid blocks. By squaring the term in parentheses in equation 2, points near the wall are assigned a very small resolution capacity, so they are solved in the local grid with an appropriate wall distance for the turbulence model.

The grids from Figure 14 are shown with the wall distance correction in Figure 19. Comparing the two figures, the effect of wall distance is clear. The small “island” above the slat is eliminated as the points in the slat grid are promoted to higher resolution capacity in the vicinity of the slat surface. The slat grid extends around the bottom of the slat and transitions to the airfoil body grid at the gap. Similarly, for the transition to the flap grid on the lower surface, the airfoil body grid extends to the gap and the flap grid is used for points through the trailing edge of the flap without the “island” of airfoil body grid in Figure 14. The airfoil body grid is very dense along its centerline behind the trailing edge of the airfoil body, but the transition from the airfoil body grid to the flap grid on the upper surface of the flap is still smooth, with enough points retained in the flap grid near the surface to resolve the boundary layer for this high Reynolds number flow. This is a particularly challenging case because both bodies are in such close proximity and the upper surface of the flap is so near a dense region of the airfoil body grid.




Figure 19: AGARD A2 airfoil grids with flap and slat fully retracted after domain connectivity operation with wall distance correction. The thick lines represent the boundaries of the grids before blanking.

Note that the correction represented in equation 2 is intended to be a proof of concept only and has not been evaluated beyond the AGARD A2 test case. It was investigated to determine if the cell and point resolution capacities could be modified in a simple way to obtain computationally better selection of grid points. As the code matures and the base of users and applications increases, a more appropriate algorithm may be developed.

5.0 Summary and Conclusions

The potential of the multi-solver paradigm and automated domain connectivity in resolving flow physics in unsteady problems which involve moving bodies has been demonstrated. The domain connectivity module (PUNDIT) shows encouraging trends in obtaining optimal domain connectivity, improving interpolation accuracy and automated parallel execution. The improvements that PUNDIT brings over the existing domain connectivity technology can be categorized into three areas:

1. *Automation*: There is no user input or intervention required for PUNDIT, which makes it completely automated and shows potential for large productivity gains. Existing technologies often require high-levels of user expertise because of the need to perform explicit hole cutting.
2. *Scalability*: None of the existing technologies provides scalable execution for all the domain connectivity procedures. We have demonstrated linear-scalability of execution if at least 1 million points are available per processor. However, this is

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

still quite short of meeting future production demands of highly-scalable execution. The main reason for the drop-off in scalability is because of solver-based mesh-partitioning, which is non-optimal for domain connectivity. This is because only a subset of processors will have intensive domain connectivity operations as the grid system becomes more finely partitioned, leaving the overall solution procedure unbalanced. Algorithms for redistribution of search load on the fly are envisioned to mitigate the load imbalance.

3. *Arbitrary element types and relative motion capability*: PUNDIT supports multiple solvers (both intra-code and inter-code), multiple mesh types and performs all the domain connectivity operations in a distributed computing platform for moving body problems. Existing technologies do support many of these features but none of them support all of the features together.

Following are the summaries of the specific observations that pertain to the various application test cases that were studied:

1. The NACA0015 test case shows the capability of adaptive meshes and high-order solvers to preserve the coherence of vortex structures. Comparison with experimental data shows underprediction of peak swirl velocity magnitudes. However these are attributed to the lift deficiency observed from the pressure distributions. Comparison with experimental data is likely to improve with the inclusion of wind-tunnel wall modeling.
2. The AGARD A2 airfoil-flap-slat test case shows the ability of PUNDIT to obtain optimal domain connectivity and facilitate smooth solution interpolation. Excellent correlation obtained with experimental data for the pressure distributions provides validation for the methodology.
3. Studies for the TRAM rotor show the capability of automated mesh adaptation and domain connectivity solution for a rotorcraft problem in a parallel processing environment.
4. Application to the aeroelastic simulation of a UH-60A helicopter further illustrates the capability of dynamic domain connectivity management in a parallel processing environment with reasonable cost.
5. Promoting the resolution capacity of grid points near solid walls can be used to obtain a more computationally effective grid. For the AGARD A2 airfoil test case, application of the wall boundary correction resulted in additional points being retained in the grid containing the solid wall.

6.0 Acknowledgments

Development was performed at the HPC Institute for Advanced Rotorcraft Modeling and Simulation (HIARMS) located at the US Army Aeroflightdynamics Directorate at Moffett Field, CA, which is supported by the Department of Defense's High Performance Computing Modernization Office (HPCMO). The authors would like to acknowledge Dr. Robert Meakin for many of the original ideas and critical concepts that have been used in the development of the domain connectivity module. We also gratefully acknowledge the contributions of Dr. Venke Sankaran in reviewing the manuscript, Dr. William Chan for providing test cases and insights, Dr. Buvana Jayaraman for verifying the robustness of the simulation capability and Dr. Beatrice Roget for discussions on search algorithms. We thank Dr. Roger Strawn and Dr. Christopher Atwood for providing the leadership and encouragement. In addition, we are grateful to Dr. Thomas Pulliam for the development of the higher-order Cartesian-mesh solver (ARC3DC). We also thank Dr. Dimitri Mavriplis and Dr. Zhi Yang at the University of Wyoming for providing the NSU3D code. Finally, we are thankful to Dr. Edwin van Der Weide and Prof. Juan Alonso at Stanford University for the use of the CHIMPS interpolation package, Dr. Ralph Noack at Applied Research Laboratory for the usage of SUGGAR module and Yikloon Lee at NAVAIR for the usage of the NAVAIR-IHC code.

References

- [1] Andrew Wissink, Jayanarayanan Sitaraman, Dimitri Mavriplis, Thomas Pulliam, and Venkateswaran Sankaran. A python-based infrastructure for overset cfd with adaptive cartesian grids. In *46th AIAA Aerospace Sciences Meeting*, number AIAA-2008-0927, Washington, DC, January 2006. AIAA. Reno, Nevada.
- [2] Jayanarayanan Sitaraman, A. Katz, Buvana Jayaraman, Andrew Wissink, and Venkateswaran Sankaran. Evaluation of a multi-solver paradigm for cfd using overset unstructured and structured adaptive cartesian grids. In *48th AIAA Aerospace Sciences Meeting*, number AIAA-2008-0660, Washington, DC, January 2008. AIAA. Reno, Nevada.

PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE AND MOVING GRIDS WITH MULTIPLE SOLVERS

- [3] Zhi Yang and Dimitri Mavriplis. Higher-order time integration schemes for aeroelastic applications on unstructured meshes. In *44th AIAA Aerospace Sciences Meeting*, number AIAA-2006-0441, Washington, DC, January 2006. AIAA. Reno, Nevada.
- [4] Jayanarayanan Sitaraman and James D. Baeder. Evaluation of the wake prediction methodologies used in cfd based rotor airload computations. In *24th AIAA Conference on Applied Aerodynamics*, Washington, DC, June 2006. AIAA. San Francisco, CA.
- [5] Stuart E. Rogers, Norman E. Suhs, and William E. Dietz. Pegasus 5: An automated preprocessor for overset-grid computational fluid dynamics. *AIAA Journal*, 41(6), June 2003.
- [6] Robert L. Meakin. Object x-rays for cutting holes in composite overset structured grids. In *15th AIAA Computational Fluid Dynamics Conference*, number AIAA Paper 2001-2537, Washington, DC, June 2001. AIAA.
- [7] Pieter G. Buning et al. *OVERFLOW Users Manual*. NASA Langley Research Center, July 2003.
- [8] Ralph W. Noack. SUGGAR: A general capability for moving body overset grid assembly. In *17th AIAA Computational Fluid Dynamics Conference*, Washington, DC, June 2005. AIAA. Toronto, Ontario, Canada.
- [9] Juan J. Alonso, S. Hahn, F. Ham, M. Herrmann, G. Iaccarino, G. Kalitzin, Patrick LeGresley, K. Mattsson, G. Medic, P. Moin, H. Pitsch, J. Schlüter, M. Svärd, Edwin Van der Weide, D. You, and X. Wu. Chimps: A high-performance scalable module for multi-physics simulations. In *42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, number AIAA-2006-5274, Washington, DC, July 2006. AIAA. Sacramento, CA.
- [10] D. M. Belk and R. C. Maple. Automated assembly of structured grids for moving body problems. In *12th AIAA Computational Fluid Dynamics Conference, Part 1*, number AIAA 95-1680, pages 381–390, Washington, DC, June 1995. AIAA. San Diego, CA.
- [11] Z. J. Wang, V. Parthasarathy, and N. Hariharan. A fully automated chimera methodology for multiple moving body problems. In *36th AIAA Aerospace Sciences Meeting*, number AIAA 1998-217, Washington, DC, January 1998. AIAA. Reno, Nevada.
- [12] William D. Henshaw David L. Brown and Daniel J. Quinlan. Overture: Object-oriented tools for overset grid applications. In *17th AIAA Conference on Applied Aerodynamics*, number AIAA 99-3130, Washington, DC, June 1999. AIAA. Norfolk, VA.
- [13] Yiklooon Lee and James D. Baeder. Implicit hole cutting—a new approach to overset grid connectivity. In *16th AIAA Computational Fluid Dynamics Conference*, Washington, DC, June 2003. AIAA. Orlando, Florida.
- [14] Robert L. Meakin, Andrew M. Wissink, William C. Chan, S. A. Pandya, and Jayanarayanan Sitaraman. On strand grids for complex flows. In *18th AIAA Computational Fluid Dynamics Conference*, number AIAA 2007-3834, Washington, DC, June 2007. Miami, Florida.
- [15] Richard D. Hornung, Andrew M. Wissink, and Scott R. Kohn. Managing complex data and geometry in parallel structured amr applications. *Engineering with Computers*, 22(3–4):181–195, December 2006. See www.llnl.gov/casc/samrai.
- [16] Juan J. Alonso, Patrick LeGresley, and Edwin Van Der Weide. pymdo: A framework for high-fidelity multi-disciplinary optimization. In *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, number AIAA-2004-4480, Washington, DC, August 2004. AIAA. Albany, New York.
- [17] Thomas H. Pulliam. Euler and thin-layer navier-stokes codes: Arc2d, and arc3d. In *Computational Fluid Dynamics Users Workshop*, March 1984. Tullahoma, Tennessee.
- [18] Richard D. Hornung and Scott R. Kohn. Managing application complexity in the SAMRAI object-oriented framework. *Concurrency and Computation: Practise and Experience*, 14:347–368, 2002.
- [19] M. J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82:65–84, 1989.

**PARALLEL UNSTEADY OVERSET MESH METHODOLOGY FOR ADAPTIVE
AND MOVING GRIDS WITH MULTIPLE SOLVERS**

-
- [20] Jayanarayanan Sitaraman, Matthew W. Floros, Andrew M. Wissink, and Mark Potsdam. Parallel unsteady overset mesh methodology for a multi-solver paradigm with adaptive cartesian grids. In *26th AIAA Applied Aerodynamics Conference*, Washington, DC, August 2008. AIAA.
- [21] S. J. Plimpton, B. Hendrickson, and J. R. Stewart. A parallel rendezvous algorithm for interpolation between multiple grids. *Journal of Parallel Distributed Computing*, 64:266–276, 2004.
- [22] Kenneth W. McAlister and R. K. Takahashi. NACA 0015 wing pressure and trailing vortex measurements. NASA TP 3151, Ames Research Center, November 1991.
- [23] A. Elsenaar et al. A selection of experimental test cases for the validation of cfd codes volume 1. AGARD AR 303, Advisory Group for Aerospace Research & Development (AGARD), 7 Rue Ancelle, 92200 Neuilly-Sur-Seine, France, August 1994.
- [24] Martin Burt. Summaries of the test cases. In *A Selection of Experimental Test Cases for the Validation of CFD Codes Volume 1* [23], chapter 5, pages 55–133.
- [25] Advisory Group for Aerospace Research & Development (AGARD). A selection of experimental test cases for the validation of cfd codes supplement: Datasets a–e. Electronic Data, August 1994. AGARD-AR-303-SUPPL.
- [26] I. R. M. Moir. Measurements on a two-dimensional aerofoil with high-lift devices. In *A Selection of Experimental Test Cases for the Validation of CFD Codes Volume 2*, pages A2–1–A2–12. Advisory Group for Aerospace Research & Development (AGARD), 7 Rue Ancelle, 92200 Neuilly-Sur-Seine, France, August 1994.
- [27] Mark Potsdam, Hyeonsoo Yeo, and Wayne Johnson. Rotor airloads prediction using loose aerodynamic/structural coupling. *Journal of Aircraft*, 43(3):732–742, 2006.
- [28] Robert M. Kufeld, Dwight L. Balough, Jeffrey L. Cross, Karen F. Studebaker, Christopher D. Jennison, and William G. Bousman. Flight testing the UH-60A airloads aircraft. In *American Helicopter Society 50th Annual Forum Proceedings*, pages 557–577, Alexandria, VA, May 1994. AHS International. Washington, DC.